

Claims

What is claimed is:

- 5 1. A method of processing a network connection in a computer system,
comprising:
- establishing the network connection by a first network protocol stack;
 determining whether to offload the network connection from the first
network protocol stack to a second network protocol stack; and
- 10 transferring the network connection from the first network protocol
stack to the second network protocol stack when it is determined to offload the
network connection from the first network protocol stack to the second
network protocol stack.
- 15 2. The method as recited in claim 1, further comprising:
 sharing state information associated with the network connection
between the first network protocol stack and the second network protocol
stack.
- 20 3. The method as recited in claim 1, wherein determining whether to
offload the network connection is performed by an operating system kernel of
the computer system.

4. The method as recited in claim 3, wherein determining whether to offload the network connection is performed by a socket layer of the operating system kernel.
5. The method as recited in claim 1, wherein determining whether to offload the network connection is performed by the first network protocol stack.
6. The method as recited in claim 1, wherein the first network protocol stack is implemented in software and the second network protocol stack is implemented in hardware.
7. The method as recited in claim 1, wherein the first network protocol stack is implemented in at least one of software and hardware, and the second network protocol stack is implemented in at least one of software and hardware.
8. The method as recited in claim 6, wherein the hardware is a TOE capable NIC.

9. The method as recited in claim 1, wherein the second network protocol stack is capable of determining whether to offload the network connection back to the first network protocol stack.

5 10. The method as recited in claim 9, further comprising:
receiving an indicator from the second network protocol stack or a driver associated with the second network protocol stack, the indicator indicating a request to transfer the network connection back to the first network protocol stack.

10

11. The method as recited in claim 10, further comprising:
obtaining state information for the network connection from the second network protocol stack or the driver associated with the second network protocol stack when the indicator is received; and

15 handling the network connection by the first network protocol stack using the obtained state information.

12. The method as recited in claim 11, wherein obtaining state information is performed by a TCP layer of the first network protocol stack.

20

13. The method as recited in claim 11, further comprising:
obtaining at least one of unsent and undelivered data by the first network protocol stack from the second network protocol stack or a driver

associated with the second network protocol stack, thereby enabling the first network protocol stack to process the unsent or undelivered data.

14. The method as recited in claim 9, wherein the first network protocol
5 stack is implemented in software and the second network protocol stack is implemented in hardware.

15. The method as recited in claim 14, wherein the hardware is a TOE
capable NIC.
10

16. The method as recited in claim 15, wherein the network connection is a TCP connection.

17. The method as recited in claim 9, further comprising:
15 handling the network connection by the first network protocol stack when the network connection is offloaded back to the first network protocol stack from the second network protocol stack.

18. The method as recited in claim 9, further comprising:
20 handling the network connection by the first network protocol stack until it is determined to offload the network connection to the second network protocol stack.

19. The method as recited in claim 1, further comprising:

handling the network connection by the first network protocol stack until it is determined to offload the network connection to the second network protocol stack.

5

20. The method as recited in claim 1, further comprising:

providing state information associated with the first network protocol stack to the second network protocol stack when it is determined to offload the network connection from the first network protocol stack to the second

10 network protocol stack.

21. The method as recited in claim 1, further comprising:

establishing a mapping between a first set of state information for the network connection maintained by or associated with the first network

15 protocol stack and a second set of state information for the network connection maintained by or associated with the second network protocol stack.

22. The method as recited in claim 1, wherein transferring the network

20 connection from the first network protocol stack to the second network protocol stack comprises:

exchanging state information for the network connection between the first network protocol stack and the second network protocol stack.

23. The method as recited in claim 22, wherein exchanging state information comprises:

exchanging a first identifier for the network connection maintained by the first network protocol stack with a second identifier for the network connection maintained by the second network protocol stack.

24. The method as recited in claim 22, wherein the state information comprises IP addresses and ports for a client and server of the network connection, and at least one of send and receive sequence numbers of one or more packets for the network connection.

25. The method as recited in claim 24, wherein the state information further comprises:

a round trip estimate.

26. The method as recited in claim 25, wherein the state information further comprises:

a congestion window and slow start information.

27. The method as recited in claim 1, wherein transferring the network connection from the first network protocol stack to the second network protocol stack comprises:

initiating the transfer of the network connection by a socket layer of the computer system.

28. The method as recited in claim 1, wherein upon transferring the network connection from the first network protocol stack to the second network protocol stack, the method further comprising:

at least one of sending one or more packets by the second network protocol stack to the socket layer and receiving one or more packets by the second network protocol stack from the socket layer.

29. A method of processing a network connection in a computer system, comprising:

establishing the network connection by an operating system of the computer system;

determining whether to offload the network connection from the operating system to a network interface card; and

transferring the network connection from the operating system to the network interface card when it is determined to offload the network connection from the operating system to the network interface card.

20

30. An apparatus for processing a network connection in a computer system, comprising:

means for establishing the network connection by a first network protocol stack;

means for determining whether to offload the network connection from the first network protocol stack to a second network protocol stack; and

means for transferring the network connection from the first network protocol stack to the second network protocol stack when it is determined to

5 offload the network connection from the first network protocol stack to the second network protocol stack.

31. A computer-readable medium storing thereon computer-readable instructions for processing a network connection in a computer system,

10 comprising:

instructions for establishing the network connection by a first network protocol stack;

instructions for determining whether to offload the network connection from the first network protocol stack to a second network protocol stack; and

15 instructions for transferring the network connection from the first network protocol stack to the second network protocol stack when it is determined to offload the network connection from the first network protocol stack to the second network protocol stack.

20

32. A network device adapted for processing a network connection, comprising:

an operating system including a first network protocol stack; and
a second network protocol stack coupled to the first network protocol
stack, the operating system being configured for determining whether to
offload a network connection to the second network protocol stack and
5 transferring the network connection from the first network protocol stack to
the second network protocol stack when it determines that it will offload the
network connection to the second network protocol stack.

10 33. The network device as recited in claim 32, wherein the first network
protocol stack is a TCP/IP stack and the second network protocol stack is a
TCP/IP stack.

15 34. The network device as recited in claim 32, wherein the first network
protocol stack is implemented in software and the second network protocol
stack is implemented in hardware.

35. The network device as recited in claim 34, wherein the hardware is a
TOE capable NIC.

20

36. The network device as recited in claim 32, wherein the second network
protocol stack is capable of determining whether to offload the network
connection back to the first network protocol stack.

37. The network device as recited in claim 36, wherein the second network protocol stack or a driver associated with the second network protocol stack sends an indicator when it requests to transfer the network connection back to
5 the first network protocol stack.

38. The network device as recited in claim 32, wherein the first network protocol stack is adapted for obtaining state information for the network
10 connection from the second network protocol stack or a driver associated with the second network protocol stack when an indicator is received, thereby enabling the first network protocol stack to handle the network connection using the obtained state information.

15 39. The network device as recited in claim 38, wherein a TCP layer of the first network protocol stack is adapted for obtaining the state information for the network connection from the second network protocol stack or the driver associated with the second network protocol stack.

20 40. The network device as recited in claim 38, wherein the first network protocol stack is further adapted for obtaining at least one of unsent and undelivered data from the second network protocol stack, thereby enabling the first network protocol stack to process the unsent or undelivered data.

41. The network device as recited in claim 38, wherein the indicator is received from the second network protocol stack or a driver associated with the second network protocol stack.

5

42. The network device as recited in claim 36, wherein the first network protocol stack is implemented in software and the second network protocol stack is implemented in hardware.

10 43. The network device as recited in claim 42, wherein the hardware is a TOE capable NIC.

44. The network device as recited in claim 43, wherein the network connection is a TCP connection.

15

45. The network device as recited in claim 36, wherein the first network protocol stack is capable of handling the network connection when the network connection is offloaded back to the first network protocol stack from
20 the second network protocol stack.

46. The network device as recited in claim 45, wherein the first network protocol stack handles the network connection until it is determined by the

operating system to offload the network connection to the second network protocol stack.

47. The network device as recited in claim 32, wherein the first network
5 protocol stack handles the network connection until it is determined by the
operating system to offload the network connection to the second network
protocol stack.

10 48. The network device as recited in claim 32, wherein the operating
system is configured to establish a mapping between a first set of state
information for the network connection associated with the first network
protocol stack and a second set of state information for the network
connection associated with the second network protocol stack.

15

49. The network as recited in claim 32, wherein the operating system is
configured for providing state information for the network connection to the
second network protocol stack when the first network protocol stack decides
to offload the network connection from the first network protocol stack to the
20 second network protocol stack.

50. The network device as recited in claim 32, wherein the operating
system is configured to at least one of provide state information associated

with the first network protocol stack and obtain state information associated with the second network protocol stack.

51. The network device as recited in claim 50, wherein the state
5 information comprises an identifier for the network connection.

52. The network device as recited in claim 50, wherein the state
information comprises IP addresses and ports for a client and server of the
connection, and at least one of send and receive sequence numbers of one or
10 more packets for the connection.

53. The network device as recited in claim 52, wherein the state
information further comprises:
a round trip estimate.
15

54. The network device as recited in claim 53, wherein the state
information further comprises:
a congestion window and slow start information.

20 55. The network device as recited in claim 32, wherein the transfer of the
network connection from the first network protocol stack to the second
network protocol stack is initiated by a socket layer of the first network
protocol stack.

56. The network device as recited in claim 32, wherein upon transferring the network connection from the first network protocol stack to the second network protocol stack, the second network protocol stack is in
5 communication with a socket layer of the first network protocol stack, thereby enabling data to be sent by the second network protocol stack to the socket layer and enabling data to be received by the second network protocol stack from the socket layer.

10